

## THE DEEP SPACE 1 EXTENDED MISSION

Marc D. Rayman and Philip Varghese  
Jet Propulsion Laboratory  
California Institute of Technology  
4800 Oak Grove Dr.  
Pasadena, CA 91109 USA

The primary mission of Deep Space 1 (DS1), the first flight of the New Millennium program, completed successfully in September 1999, having exceeded its objectives of testing new, high-risk technologies important for future space and Earth science missions. DS1 is now in its extended mission, with plans to take advantage of the advanced technologies, including solar electric propulsion, to conduct an encounter with comet 19P/Borrelly in September 2001. During the extended mission, the spacecraft's commercial star tracker failed; this critical loss prevented the spacecraft from achieving three-axis attitude control or knowledge. A two-phase approach to recovering the mission was undertaken. The first involved devising a new method of pointing the high-gain antenna to Earth using the radio signal received at the Deep Space Network as an indicator of spacecraft attitude. The second was the development of new flight software that allowed the spacecraft to return to three-axis operation without substantial ground assistance. The principal new feature of this software is the use of the science camera as an attitude sensor. The differences between the science camera and the star tracker have important implications not only for the design of the new software but also for the methods of operating the spacecraft and conducting the mission. The ambitious rescue was fully successful, and the extended mission is back on track.

### INTRODUCTION

Deep Space 1 (DS1) was launched on October 24, 1998 as the first mission in NASA's New Millennium program (NMP). NMP is designed to accelerate the implementation of ambitious space science missions by developing and testing some of the high-risk, high-benefit technologies they need. NMP conducts deep space and Earth orbiting missions focused on the validation of these technologies.<sup>1</sup> Thus, the 11-month primary mission of DS1 was devoted to the testing and evaluation of 12 technologies selected on the bases of their importance to future space and Earth science programs, the significant advancements they offer over current state-of-the-art, the high risk they present to the first user, and the need for in-flight testing to reduce that risk.

In addition to its technical objectives, DS1 was intended to probe the limits of rapid development for deep-space missions. The initial study of DS1 was undertaken only 39 months before launch, an unprecedentedly short time for

a NASA deep-space mission in the modern era. At the time the preliminary concept study was initiated, the only definition of the project was that it would validate solar electric propulsion and other unidentified technologies in deep space and that launch would occur sometime in 1998. The level-1 requirements and goals were formulated 26 months prior to launch.

The advanced technologies tested during the primary mission are:

- solar electric propulsion
- solar concentrator arrays
- miniature integrated ion and electron spectrometer
- miniature integrated camera and imaging spectrometer
- autonomous optical navigation
- small deep-space transponder
- K<sub>a</sub>-band solid state power amplifier
- beacon monitor operations
- autonomous remote agent
- low-power electronics
- power actuation and switching module
- multifunctional structure

The technologies and their performances during the primary mission have already been described.<sup>2,3</sup> All but the last four technologies listed above play some role in the extended mission. Some of them, including the solar arrays, transponder, and the autonomous navigation system (the real-time part of which provides positions of the Sun, Earth, and other bodies to the attitude control system), are critical to basic spacecraft functioning.

All of the technology testing was completed by July 1999. Following the end of the testing, DS1 conducted a bonus encounter with asteroid (9969) Braille on July 29, 1999 (described below). The primary mission ended on September 18, 1999, having met or exceeded all of the mission success criteria. The total cost, including launch but excluding the development of some of the technologies, from the commencement of the project through the end of the mission was \$149.7 M (in real-year dollars). More detailed background on the Deep Space 1 project, the mission, and the spacecraft has been presented elsewhere.<sup>1,2,4</sup>

### BRAILLE ENCOUNTER

Because the encounter was lower priority than the technology validation and the operations team was small, detailed planning activities for the encounter with Braille began after the technology testing entered its final phase. A rehearsal covering the final 7 hours of the encounter was conducted very successfully on the spacecraft on July 13. It included the autonomous design and execution of trajectory correction maneuvers (TCMs) and even a system on board to intercept images taken for optical navigation, “paint” a synthesized asteroid in them, and reinsert them into the data transfer and processing path. This allowed all characteristics of the camera performance (including flaws but excluding its actual response to the real asteroid) to be fully expressed in the rehearsal.

DS1 is capable of performing TCMs with either the ion propulsion system (IPS) or the hydrazine-based reaction control system (RCS), which is normally used only for attitude control. IPS TCMs could take up to one day to complete but enable the spacecraft to modify its trajectory while using less propellant. Initially TCMs were conducted with the IPS, but beginning at about two days before closest approach (CA), space-

craft activity increased as the autonomous navigation system (AutoNav) increased its frequency of optical navigation imaging, so DS1 executed the TCMs then with the RCS, which accomplishes the corrections much more quickly.

AutoNav used the visible charge coupled device (CCD) in the miniature integrated camera/spectrometer (MICAS) to obtain optical data with which to update its trajectory knowledge throughout most of the primary mission. On approach to the asteroid, it used these data, combined with other data available on board,<sup>2,3</sup> to determine its trajectory. RCS TCM opportunities were placed at about CA – 50 hours (h), CA – 33 h, CA – 18 h, CA – 12 h, and CA – 6 h, although it was not expected that all of them would be used. Rather, they were windows during which AutoNav was invited to execute a TCM if it deemed one to be necessary.

Braille was so dim that the initial firm detection of it was with ground analysis of MICAS images acquired at CA – 40 h and analyzed in time for the TCM at CA – 33 h. The asteroid turned out to be about 430 km from the best prediction of its location. Because of the large targeting error that resulted from this  $1.6\sigma$  ephemeris error, a prompt TCM was deemed essential. The design used the AutoNav software running on the ground (because the spacecraft had not yet detected the target) with improved picture processing that was unavailable to the spacecraft. The TCM was uplinked and executed by AutoNav as if it had designed the maneuver itself.

AutoNav did lock on to the asteroid with its images acquired at CA – 17 h. At the end of the imaging session however, at approximately CA - 16 h, a software bug caused the spacecraft to enter one of its safe states. This prevented AutoNav from planning and conducting a TCM at CA - 12 h based on the CA – 17 h data. While the operations team was rapidly identifying the cause of the safing and bringing the spacecraft back to normal operational configuration, 3 of the optical navigation images acquired shortly before the safing event were recovered and downlinked. Those images were used to design the CA - 6 h TCM using the AutoNav software running on the ground. The TCM was uplinked in time for the encounter sequences to resume at CA - 6.5 h, with AutoNav executing the TCM. Because only 3 images were available for the

ground to design the TCM (AutoNav would have used 16), the accuracy was not as high as AutoNav would have been able to achieve had the safing not interfered. As a result, the spacecraft's closest approach to Braille was  $28.3 \pm 1.5$  km rather than the originally planned 15 km. At about 2 km in diameter, Braille is the smallest solar system body ever targeted for an encounter; for this case,  $v_{\infty} = 15.5$  km/s.

The complex encounter sequences executed as intended. AutoNav successfully tracked the asteroid to CA - 70 minutes (m) using the visible CCD in MICAS. Because of prelaunch predictions that the CCD would not be useable when a bright body spanned many pixels, at CA - 27 m AutoNav was commanded to switch to a mode in which it used the active pixel sensor (APS) in MICAS for computing its updates of the asteroid location. Even with conservatively selected MICAS integration times and detection thresholds, the asteroid never reached threshold in any of the 23 images taken for AutoNav. Indeed, the actual signal:noise probably never even reached unity. This was later understood to be the result of the asteroid being dimmer than the worst case predictions available and an unknown nonlinearity in the APS response under weak illumination. As a result, the spacecraft pointing was not adequate during the final 150 seconds before CA to acquire data with MICAS. All data from the plasma experiment for planetary exploration (PEPE) and the suite of IPS diagnostic sensors (IDS), reprogrammed for encounter science data acquisition, were collected as planned, with AutoNav initiating the final four sequences based on its estimated time of closest approach from the CCD data gathered prior to CA - 70 m. The spacecraft was commanded not to track the asteroid through CA, in order to reduce RCS activity so that PEPE and IDS would have as clean an environment as possible. Shortly after closest approach, the spacecraft turned to point MICAS at the asteroid. AutoNav used its estimate of the asteroid location from the CA - 70 m CCD observations, and MICAS acquired visible and infrared data at CA + 15 m.

### MISSION PROFILE

When DS1 was launched, the plan for the primary mission incorporated the encounter with Braille. The mission design maintained an option for an extended mission encounter with comet 19P/Borrelly in September 2001. After

launch the mission progressed so well, with the critical technologies exhibiting excellent performance, that the proposal for the extended mission was augmented to include an encounter with comet 107P/Wilson-Harrington. The extension to the mission was approved by NASA in August, 1999. That extended mission was described by Rayman *et al.*<sup>2</sup>

Thrusting with the IPS during the primary mission was designed to allow extensive testing of the technology, and it placed the spacecraft on the trajectory to Braille and the extended mission targets. Although the request for the extended mission had not yet been granted, thrusting resumed only about 36 hours after the closest approach to Braille so that if the proposal were approved, the spacecraft would already be on course for the comets.

A typical week of IPS thrusting began with AutoNav commanding the spacecraft to turn to point the ion engine in the direction required for thrusting. AutoNav then started the IPS and throughout the week updated the thrust direction and throttle level. After about 150 hours, the IPS was turned off and AutoNav commanded the attitude control system (ACS) and MICAS to collect CCD images of selected asteroids and background stars for its use in on-board orbit determination. The collection of such images could last for up to 4 hours, at the end of which AutoNav pointed the unarticulated high-gain antenna (HGA) to Earth for the weekly track by the Deep Space Network (DSN). At the end of the DSN session, AutoNav took the spacecraft back to the thrust attitude. The workload for the operations team was significantly less than it would have been without AutoNav.

Trajectories that use solar electric propulsion and are optimized to maximize the neutral mass (defined to be all flight mass except IPS propellant) delivered to the target typically have periods in which coasting is better than thrusting. Because no encounters were required for the DS1 primary mission, the timing of thrust and coast periods was determined in large part by the technology experiments; some tests required thrusting, and some required coasting. Only encounter targets that allowed coast periods at times that matched the needs of the extensive technology testing program during the primary mission were considered. Some coast periods in the thrust profile were optimal and

others were inserted to allow special spacecraft activities or to provide buffer against periods of unexpected loss of thrust.

Following nearly three months of thrusting after the Braille encounter (the longest thrust arc of the mission at that time), a coast period began on October 20, 1999. At that point, the IPS had consumed 21.6 kg of the 81.5 kg of xenon on board at launch, imparted 1.32 km/s to DS1, and completed 3571 hours of operation.

During the coast period, planned to last until the middle of December 1999, extensive new calibrations of all channels in MICAS were conducted. In addition, MICAS acquired 48 infrared spectra of Mars, covering nearly two full rotations of the planet, from a range of 55 million km. With three spectra collected every three hours for 48 hours, the data allow resolution of  $45^\circ$  in longitude. These data are considered to be the highest quality spectra of Mars ever collected in the range of  $1.3\ \mu\text{m}$  to  $1.9\ \mu\text{m}$ . Spectral features have been detected that may indicate the presence of previously unrecognized surface minerals.<sup>5</sup>

#### STELLAR REFERENCE UNIT FAILURE

On November 11, 1999, after all the Mars spectra were acquired but before the next scheduled DSN track, during which the data would be returned, the spacecraft's stellar reference unit (SRU) stopped reporting attitude data to the spacecraft computer. By recognizing star patterns, this commercial unit produced a quaternion and thus was able to provide the complete three-axis attitude. The SRU was one of three attitude sensor types; the spacecraft also carries one laser gyro for each axis and a Sun sensor assembly (SSA), with  $128^\circ$  full-angle field of view, which was used principally for safe modes. Attitude is normally controlled using the RCS; when the IPS is thrusting, ACS controls two axes by moving the ion thruster through a range of  $\pm 5^\circ$ . This thrust vector control (TVC) system using the IPS substantially reduces hydrazine expenditure.

The SRU had exhibited intermittent problems since shortly after launch. Diagnostic activities on board and laboratory tests conducted jointly between the vendor and JPL had not yielded an explanation for the occasional interruptions in its reporting of attitude. The

longest outage had been 28 minutes.

When the SRU exhibited problems on November 11, the spacecraft's fault protection system power cycled it two times, neither one of which cleared the problem, before finally declaring a celestial inertial reference loss (CIRL). CIRL leads to the spacecraft entering a safe state known as Sun standby SSA. To achieve this state, the SRU and some other devices are power cycled, non-essential devices are turned off, and ACS uses the SSA and gyros to point the spacecraft's  $+x$  axis at the Sun and rotate around the Sun-spacecraft line at 1 revolution/hour. The center of the SSA's field of view and the center of the HGA are along the  $+x$  axis. ACS also rotates the solar arrays so that they are normal to  $+x$ .

The spacecraft has three low-gain antennas (LGAs): one each aligned with  $+x$ ,  $+z$ , and  $-z$ . The HGA and LGAs all work at X-band. (There is also a  $+x$   $K_a$ -band antenna that has been used principally for technology experiments and for DSN testing but also can return telemetry to the few DSN stations equipped for  $K_a$ -band reception.) The rotation triggered by CIRL is a remnant from a very early mission phase in which the Sun-probe-Earth (SPE) angle was too large to return telemetry through the  $+x$  LGA in Sun standby SSA; the LGAs on the  $z$  axis were selected then. For the remainder of the extended mission, the SPE angle will remain less than  $45^\circ$ , so the  $+x$  LGA will always be used in this safe state; at the time of the SRU failure, the SPE angle was  $38^\circ$ . Two-way communications required two DSN stations. At a geocentric range of 1.6 astronomical units (AU), the 34-m stations of the DSN were below threshold for telemetry, although they were capable of commanding the spacecraft. The 70-m stations could support low-rate downlink, but they did not have X-band uplink. (Since then, the 70-m station at Goldstone Deep Space Communications Complex has been augmented with X-band uplink, and the same upgrades for the 70-m stations at the Canberra and Madrid Deep Space Communications Complexes are scheduled to be completed in November 2000 and October 2001 respectively.)

The initial analysis of the SRU failure was severely hampered by the limited downlink rate and the sparse DSN coverage that had been scheduled. A large volume of engineering

telemetry needed to be returned to provide the complete context for the SRU's anomalous behavior. Before the return of all the data, several attempts to revive the unit were conducted, all without success. By the end of November, evidence that the SRU could not be resuscitated was accumulating.

During development and operations the SRU was considered a critical spacecraft device. Funding was not adequate however to provide for redundancy. As a result, loss of the SRU had always been considered a mission-ending failure. Given that the failure occurred after the successful conclusion of the primary mission, one of the options was to terminate the extended mission. Nevertheless, the project decided to undertake an extremely rapid and extensive recovery effort in two phases.

### PHASE 1 RECOVERY

It was clear that to conduct a thorough diagnosis of the SRU, to return the large volume of Mars data, and to pursue any further meaningful activities with the spacecraft, it would be necessary to use the HGA. Thus, the first phase of a recovery was initiated, with the objective being to point the HGA to Earth. Based on tests conducted in the DS1 testbed at JPL and on the spacecraft, an experimental procedure was developed and executed successfully on January 14, 2000.

With only gyros and the SSA, ACS has knowledge of the Sun location but no other celestial reference. To point the HGA to Earth, the first step was to command the spacecraft to offset the Sun from the center of the SSA by the SPE angle. On January 14, that angle was 34.3°. Although the direction of offset could be specified in spacecraft body coordinates, the relationship of that direction to inertial space was unknown. Once the offset was achieved, the rotation rate around the Sun-spacecraft line (now 34.3° from the  $+x$  axis) was commanded to 1 revolution/45 minutes. The spacecraft transmitted an unmodulated carrier through its HGA as it coned around the Sun.

In Sun standby SSA, the solar arrays are caged so that the plane of the panels is normal to  $+x$ . Because the arrays use cylindrical concentrator lenses, they are very sensitive to pointing in one axis; thus, extra commands were

included in the offset turn sequence to rotate the arrays back by the SPE angle so that they would remain orthogonal to the Sun-spacecraft line.

The DSN observed the X-band signal as the spacecraft eventually swept past Earth, revealing the unknown phase of the rotation. The maximum of the received signal power corresponded to the HGA being Earth-pointed. Because of uncertainty in gyro bias values, two maxima were used to refine the knowledge of the rate of the rotation. Initially there was also uncertainty in the time it would take ACS, still operating in its Sun standby SSA mode, to achieve steady-state rotation, so the first maximum was observed but not used for measuring the rate.

Once the phase and rate of the coning were known, the time of the next maximum was predicted. A special short uplink frequency sweep had been developed with the DSN, and, accounting for the one-way light time (13 minutes 55 seconds on January 14), the beam-width of the HGA (4° from the center to the 3 dB point), and the rotation rate, the time of the beginning of the sweep was computed. The sweep would begin at a time that would make it arrive at the spacecraft as the leading edge of the HGA moved Earth into its beam. The sweep completed, thus bringing the spacecraft receiver into frequency lock, in time to allow one command to be uplinked before the rotation would take the trailing edge of the HGA out of view of Earth.

The single command that was transmitted activated a sequence that had been uplinked earlier through the LGA. The sequence commanded the spacecraft to stop coning around the Sun. Analyses and tests had enabled predictions for the duration of the deceleration as well as the expected time past the point of  $+x$  on Earth that the command would be received (given the uplink sweep time, command radiation time, and other delays in the system). Thus, the sequence included commands for the spacecraft to rotate back far enough to account for these effects. When the next signal peak was observed, the signature of the carrier power at the DSN clearly revealed the spacecraft continuing past it and eventually backing up, ending with the HGA within 2 dB of the predicted value for optimal Earth-pointing.

Once the HGA was pointed to Earth, gyro

biases would gradually move it away. ACS would keep the Sun at the commanded offset angle from the center of the SSA, so the effect of the gyro biases was to rotate the spacecraft around the Sun-spacecraft line. Ten sequences, each commanding the spacecraft to rotate a fixed angle from  $-10^\circ$  to  $+10^\circ$  in  $2^\circ$  increments, were stored on board. As the observed carrier power at the DSN diminished to the point that a correction was deemed necessary, one of the 10 sequences would be activated by real-time command, based on how much the signal had decreased and whether it was seen to have passed through the peak before decreasing (thus indicating the sign of the correction that was needed).

Real-time commands were used to select the uplink and downlink rates, based on how close the HGA was to Earth-point (as inferred from the signal strength at the DSN and, in later activities, telemetered measurements of uplink carrier power received at the spacecraft). At the end of a DSN session, the spacecraft was commanded to point  $+x$  to the Sun again. If it were left at the SPE offset angle, gyro biases during the gap in DSN coverage (typically 1 week) would have been sufficient to make the spacecraft attitude and, therefore, the LGA direction, unpredictable. With  $+x$  pointed at the Sun, the uncertainty in phase around the Sun-spacecraft line from the gyro bias was irrelevant for LGA communications and for the initiation of subsequent repointing activities.

Once the HGA was Earth-pointed, the first priority was return of the Mars spectra. With DS1's future being in grave doubt, given the inability to revive the SRU, it was considered most important to return the science data that had already been acquired.

The capability to point the HGA to Earth allowed a more complete investigation into the state of the SRU to begin. Extensive diagnostic activities were conducted that simply would have been far too data intensive through the LGA. As JPL worked with the unit's manufacturer, it became evident that its failure was indeed permanent, as had been suspected. Despite a significant effort on behalf of subsequent users, however, the failure mechanism could not be established.

The pointing procedure proved extremely successful and productive, and it was used

many times from January through June. When the HGA was Earth-pointed, attitude corrections needed to be transmitted only about once every 2 hours. Nevertheless, the procedure did consume valuable DSN time during the coning, and each activity required planning (to account for changing SPE angle and geocentric range) and diligence that was not negligible for the very small DS1 operations team. In addition, although the spacecraft could be controlled to point  $+x$  to Earth, this technique did not permit a practical way to achieve any other attitude.

## PHASE 2 RECOVERY

The permanent loss of the SRU meant that if any further worthwhile operations were to be attempted with the spacecraft, a new method of controlling the attitude would be necessary. Even returning to technology validation would be impractical without some changes. For example, further tests with the IPS would produce a small torque that ACS would not be able to counter with the ion thruster gimbal in the absence of SRU data, so it would have to use the RCS. But by the end of January, about 16 kg of hydrazine (from an initial load of 31 kg) remained on board; this was insufficient to control the attitude for any meaningful duration of IPS thrusting. Thrusting without using thrust vector control would be extremely costly.

In late January 2000, in parallel with detailed investigations into the SRU, several methods for replacing the attitude knowledge that had been provided by the SRU were considered, but the one that was selected relied upon using the visible CCD channel in MICAS to track a star for attitude reference. Differences between MICAS and the SRU made this replacement far from straightforward however. Table 1 shows some of those differences.

The new method was complicated by the presence of scattered light in MICAS. The scattered light was studied extensively as part of the mission's technology validation experiments and was well understood from in-flight testing and modeling. In many attitudes it reduced the useful field of view by about 30% from what is shown in Table 1. Furthermore, the combination of scattered light and regions of the MICAS field of view (FOV) with decreased sensitivity limit the faintest star that could be used for reliable attitude reference to  $m_v \approx 6$ .

Parameter	SRU	MICAS
Field of view	$8.8^\circ \times 8.8^\circ$	$0.69^\circ \times 0.78^\circ$
Limiting stellar magnitude	$m_v = 7.5$	$m_v \approx 10$
Output format	Quaternion	Image file
Output rate to spacecraft computer	4 Hz	0.03 Hz

Table 1. Key differences between the SRU and the visible CCD channel in MICAS. Some of these differences became even greater given implementation issues, as described in the text.

A major challenge with the use of MICAS in place of the SRU is that in an arbitrary attitude, the probability of a detectable star being in the MICAS FOV is too low. The solution chosen is to constrain the spacecraft to attitudes that satisfy one of two criteria: either one and only one preselected star of sufficient magnitude ( $m_v < 6$ ) is in the MICAS FOV or the duration at the attitude is short enough that gyros can be used. In contrast to the primary mission, the remainder of the extended mission could accommodate such a requirement, with four classes of attitudes needed: HGA on Earth, IPS thrusting to the comet, trajectory correction maneuvers, and science instrument recalibration and data acquisition at Borrelly.

The DS1 project elected to attempt a complete recovery of the spacecraft in time to provide an opportunity to conduct a comet encounter. Throughout development and operations a considerable body of work had been devoted to analyzing the trajectory DS1 was planning to follow; no low-thrust trajectory had ever been studied in as much detail. An important figure of merit for a low-thrust mission was determined to be the robustness to unexpected missed thrusting. While anomalies short enough to cause significant problems for a conventional chemical propulsion mission (such as missing a major trajectory correction maneuver) are unlikely to have much effect on a low-thrust mission, the mission may still be endangered by long-periods of missed thrusting. Several techniques were employed to build margin into DS1's trajectory, and it could accommodate periods of well over a month of lost thrust. The extensive phase 2 recovery operations however would exceed the time that

the spacecraft could miss IPS thrusting and still reach both targets. As a result, it became necessary to abandon at least one of the comets in the extended mission. The Deep Space 1 Science Team selected the original extended mission target, comet Borrelly, over comet Wilson-Harrington. To reach Borrelly, IPS thrusting had to resume by late July 2000. It was believed that with two to three more months to work on the problem, the probability of success would be significantly higher, but the opportunity to have a chance for a comet encounter led the project to pursue the more aggressive plan.

NASA and JPL approved the ambitious and very risky second phase of the recovery. Because the primary mission had already concluded successfully, the consequences of a failure were deemed low. The likelihood of success, as perceived by the DS1 project and communicated clearly and frequently to management at JPL and NASA, also was low.

Work on a new system began in February. During four months, software and operational methods were designed, developed, tested, and integrated. In addition to testbed testing, some developmental tests were conducted on the spacecraft.

Although scattered light in MICAS had been investigated in detail during the primary mission, there were no data on the signature at some attitudes that would be important during operations with the new system. Scattered light is independent of roll angle around the Sun-spacecraft line, so further measurements were possible even without three-axis knowledge or control. MICAS' boresight is parallel to the +z axis, so to verify that the scattered light models were correct in attitudes that had not been explored previously, the spacecraft was commanded to turn in the same way as before an HGA pointing session. Thus, in these tests the desired turn was accomplished by having ACS move the Sun by the desired angle from the center of the SSA. Images were collected and returned during HGA tracks. These data validated the scattered light models and yielded improved confidence in selecting appropriate parameter values for the new system.

Although MICAS was used frequently during the primary mission, both for validation

of it as a new technology and by AutoNav for its optical navigation images, it would be used much more extensively for the remainder of the extended mission. It includes no moving parts to wear out (part of the technology innovation), but to assure that it would be capable of providing reliable attitude information, in March a test began in which every 30 seconds one MICAS image was acquired and transferred to the spacecraft computer. A problem that had been suspected from limited evidence both in the testbed and the spacecraft manifested itself in this long-duration test. On very rare occasions, several data words are dropped somewhere along the way from MICAS to the board that provides the interface to the spacecraft computer. This renders the image file useless and sometimes causes the transfer of subsequent images to stop. The problem happens so rarely that it had not shown up clearly before, but with images planned essentially continuously for the rest of the mission, it had to be accommodated. It was determined that the addition of a simple command before every image request would clear this and some related problems; although the previous image would be lost if the words were missing, subsequent images could transfer normally.

The new system that was developed requires one and only one preselected bright star in the MICAS FOV (multiple reference stars could cause confusion and be difficult to track with normal deadbanding). Initialization and acquisition are discussed below, but once ACS has acquired that star, it tracks it by issuing a request to MICAS to take an image which is delivered to AutoNav for processing. AutoNav subtracts a background image to suppress some of the scattered light effects and locates all the candidate stars, some of which may be cosmic rays. Building upon the existing capability in AutoNav to process MICAS images was crucial for timely completion of the software.

The locations and integrated intensities of candidate stars are delivered to ACS. Ground commands inform ACS what the stellar magnitude is for the star to be tracked, and ACS identifies the star from among the candidates found by AutoNav. ACS includes limits on how much the observed magnitude is allowed to differ from the expected magnitude. This helps account for several effects, including certain regions in the MICAS FOV with greatly reduced

optical throughput. Further discrimination is provided by using estimates of spacecraft motion (as measured by gyros and the SSA) to predict where the star should be in each image, based on where it was observed in earlier images. ACS then incorporates the measured location of the star into its control loop. The system has protections built in to accommodate a missed picture or a picture in which the star fails to be detected.

To acquire a star, ACS is instructed to turn (relying on gyros and the SSA during the turn) to the target attitude. When it arrives at the estimated location, it collects the image to be used by AutoNav for background subtraction. It then begins a mosaic with MICAS. The mosaic size can be adjusted;  $3 \text{ FOV} \times 3 \text{ FOV}$  (with some overlap from each element to the next) is normally used. At each element of the mosaic, it acquires two images and searches for the candidate star in each (to avoid confusion from cosmic rays). If a star close enough to the desired magnitude is found, the mosaicking is terminated and ACS transitions from acquisition to tracking. If no star is observed within that range, the mosaic continues, although rejected stars are catalogued. Upon completion of the full mosaic, if a star within a broader range was observed somewhere in the mosaic, it is used. If that criterion is not satisfied, a new mosaic is begun. That mosaic can overlap the previous one or be moved to a new location by a desired angle, depending upon the values in parameters that are easily updated by ground command.

The celestial coordinates of the target star are included with its magnitude (and MICAS integration time) in the ground-generated commands; ACS assumes when it has found a star that is consistent with the observed magnitude and SSA angle that the star is the correct one.

Including a star catalog on board was considered and rejected during development. The very ambitious schedule led to a decision to limit the complexity of the on-board system and instead rely on ground tools to generate the commands to include all necessary information for each turn and subsequent acquisition.

It had been planned that new software would be loaded during the extended mission for the comet encounters. As a result, the DSN schedule already had adequate coverage in April

2000 for such a tracking-intensive activity. To allow more time for development (at the expense of less time for in-flight testing), the DSN agreed to postpone the tracking allocation. Coverage for uplinking the software began on May 30.

DS1 reloaded the flight software three times during the primary mission, in all cases primarily to enable new technology validations. Thus the process for replacing the software was well understood. The principal differences between this new software load and previous ones were that uplink rates were lower (because of greater geocentric range -- the spacecraft was 2.0 AU from Earth at this time) and each pass that did not have a hand-over from a preceding pass had to begin with the time consuming process of bringing the HGA to Earth-point. To load the 4 megabytes of software required 267 command files.

On June 8, 2000 the computer was commanded to reboot and install the new software. By this time, the updated trajectory analysis, based on new operational principles discussed below, showed that to maintain adequate margin for reaching Borrelly, thrusting should resume by about July 5. As a result, an intensive test and verification campaign was necessary as soon as the new software was running on board. But because of the fast pace of the work leading up to loading the software, in-flight tests had not been designed in detail. A rapid cycle of design, development, testbed testing, and spacecraft execution of tests of both the flight software and operational procedures was undertaken.

The new capabilities were activated and tested methodically but quickly. One feature of the new ACS is the use of all available data (from the SSA and, when tracking, MICAS images) to estimate gyro biases. (The excellent knowledge provided by the SRU had made including SSA measurements in this estimation unnecessary.) Thus, during the HGA passes even before locking to a star, the attitude stability, as revealed by the carrier power detected at the DSN, was much improved.

The first attempt to lock to a star was on June 12. The initialization discussed here represents a combination of what was executed then with the general procedure that was

developed for subsequent use in the event that the spacecraft loses its attitude knowledge. The HGA is brought to Earth-point, but now with the gyro biases estimated. The star pattern near the MICAS FOV is predicted for the HGA being pointed to Earth and the Sun being at the known offset angle in the SSA. The dominant uncertainty in this is in the ability to determine the angle between the HGA boresight and the Earth-spacecraft line from the measured signal strength at the DSN. The accuracy with which this knowledge could be fed back to the spacecraft is estimated to be 4°; one contribution to this error is the effect of the unpredictable component of the gyro biases during the round-trip light time.

A bright star or loose grouping of bright stars near the MICAS FOV is selected, and the spacecraft is commanded to turn to it and begin mosaicking. Once ACS finds a star and begins tracking it, telemetry shows the measured stellar magnitude. If that is insufficient to determine unambiguously which star is being tracked, a deep image is taken in place and downlinked. Such an image reveals fainter stars than the on-board system can detect and aids in making a positive star identification. If any ambiguity remains, ACS is commanded to stop tracking the star, a short turn is executed to a chosen location, an image is collected, and the spacecraft returns to resume tracking the star. This image shows nearby stars to confirm the attitude.

When the spacecraft is locked to the star, it can remain there as long as necessary and thus is very stable. Once the star identification is complete, if the spacecraft is not locked to the planned star, the quaternion that corresponds to the actual star is uplinked. The spacecraft then has a complete and accurate knowledge of its three-axis attitude. Experience has shown that once this is complete, the reliability of locking to other preselected stars after turns, even when turning > 50°, is extremely good.

Tests conducted included turns to new stars, methods to acquire the attitude in the event it is lost, and tuning of parameters to make the system more robust. On June 21, after a hiatus of more than 7 months, the IPS was turned on to test ACS' ability to achieve thrust vector control. All tests were completed with excellent results.

## RESUMPTION OF ROUTINE OPERATIONS

An important aspect of the recovery effort was the design of a new trajectory to reach comet Borrelly. Some of the challenges of low-thrust mission design are described by Rayman *et al.*<sup>6</sup> The new trajectory however had to satisfy a new constraint: all thrusting had to be in attitudes with a preselected bright star in the MICAS FOV. Studies showed that the trajectories that assumed continuously variable thrust directions (implemented before the SRU failure by AutoNav updating the thrust attitude every 12 hours) could be broken into a small number of discrete segments, each with thrusting in a fixed inertial direction. The 8 months of thrusting needed to reach comet Borrelly could be accomplished with as few as 3 segments, although to make the design more flexible, about 10 segments are used. Each segment uses one reference star, designated a “thruster,” for ACS to track. Trajectory control is achieved by adjusting the duration of thrusting each week and the transition time from one thruster to the next.

The project had set for itself a very aggressive, success-oriented schedule that included the resumption of thrusting on July 5. That would have allowed reasonable margin for later unexpected losses of thrust. The testing in June went so well, however, that thrusting to the comet began on June 28.

A typical week of thrusting with the new system is very similar to thrusting before the SRU failure, although AutoNav is not used for this. There was not enough time to make the major changes in AutoNav that would have been required for it to implement the new operational procedures. To begin, the spacecraft turns to the thruster, and ACS acquires and tracks it. The IPS is activated by stored sequence and thrusts throughout most of the week. Throttle levels during the week are chosen in advance and commanded from a sequence. Shortly before the next scheduled DSN pass, the spacecraft stops thrusting and turns to a star (designated an “Earth star”) that has been selected for that date such that when MICAS is pointed at it, the HGA is close to Earth-point.

Because of the high rate of hydrazine expenditure during the phase 1 HGA pointing activities and some of the tests early in phase 2,

the hydrazine margin for attitude control for the remainder of the mission is small. To reduce hydrazine consumption, the IPS thrusts when the HGA is on Earth-point, thus allowing ACS to control two axes with xenon instead of hydrazine. When Earth-pointing is close to the desired thrust attitude, the IPS is operated at a high throttle level, thereby aiding in reaching Borrelly; otherwise, it is operated at a very low throttle level, providing ACS adequate control authority but causing minimal effect on the trajectory. This scheme has significantly reduced hydrazine use. As it turns out, the thrust attitude is close to the Earth-point attitude for most of the period of deterministic thrusting.

The new system has proven to be quite successful. On only one occasion between the first lock (on June 12) and September 25 did the spacecraft lose track of a star without reacquiring it on its own. Apparently because of very high solar activity, which affected a number of spacecraft, on July 16 there were too many false star candidates in the pictures, and ACS lost track of the star. The spacecraft continued IPS thrusting in the desired attitude however, using the gyros and SSA. As it slowly drifted with the gyros, ACS mosaicked until it found a star that satisfied the criteria it had been applying. That stopped the drift, and allowed the gyro bias estimates to be updated, although the star was not the desired one. It eventually lost that star and found another one. Nevertheless, its estimates of gyro bias were accurate enough that it had not drifted far from the thruster. When the next DSN track was scheduled, on July 18, DS1 turned from its thrusting attitude, which was treated as being correctly locked on the thruster, to the expected location of the Earth star. It began with the wrong star, so it turned to an incorrect attitude and could not locate the Earth star. But it was close enough that the HGA was on Earth-point. The operations team quickly discovered the problem. Following the procedure described earlier, the spacecraft was commanded to return a deep image, and later was directed to a nearby star. The lost thrust time in this case was negligible.

The resumption of long-term thrusting has incidentally enabled DS1 to set the record for the longest operating time of any propulsion system in space. On September 25, 2000, DS1 had more than 5800 hours of operation on the IPS (which had consumed 32 kg of Xe and provided

a  $\Delta v = 1.9$  km/s). The previous record was held by NASA's Space Electric Rocket Test (SERT) II, launched in 1970 and operated in Earth orbit. SERT II accumulated 3879 hours of operation on one of its two experimental ion engines before the engine failed from an internal short.

### MISSION PLAN

Each day of thrusting now consumes about 0.1 kg of Xe and yields a  $\Delta v = 7$  m/s. To maintain margin, the trajectory that is being followed does not assume thrusting for the 28 days around solar conjunction, on November 11, 2000. Any thrusting that can be accomplished then will increase mission margins further, but flying a profile that relies on thrusting during that period is unnecessarily risky. Because telecommunications will be difficult or impossible for about 3 weeks around conjunction, any loss of thrust would not be correctable promptly.

In March 2001 new software will be loaded to provide the spacecraft with capabilities needed for the encounter with comet Borrelly 6 months later. This will include some changes in autonomous encounter pointing (e.g., using the visible CCD instead of the APS and methods to find the nucleus in the presence of the optically confusing coma) that had been designed and tested prior to the SRU failure but were not included in the June 2000 software load. During this next software load period, all IPS thrusting will be on Earth-point and will contribute to reaching the comet. In-flight encounter rehearsals will be conducted shortly after the software load and in June 2001.

The trajectory plan completes deterministic thrusting in April 2001, but well over 1 month, and in some cases (depending upon the date) up to 4 months, of lost thrust can be accommodated. The trajectory is shown in Figure 1.

TCMs will be executed as the spacecraft approaches the comet. For TCMs that last longer than the time allowed on gyros, the maneuvers will be decomposed into components whose vector sum achieves the required correction, with each component aligned with a reference star.

Closest approach to comet Borrelly will occur on September 23, 2001 at 1.36 AU from the Sun, about 10 days after the comet's

perihelion. The baseline plan is for the spacecraft to pass at 17 km/s about 2000 km from the nucleus on the Sun-nucleus line. (The spacecraft was not designed to encounter a comet, so managing the risk of dust impacts to the spacecraft will be a key criterion in the final selection of that distance.) Two of the technologies tested during the primary mission were compact science instruments, each with a broad range of measurement capabilities integrated into one small package. Infrared spectra and images can be obtained with MICAS, and PEPE can sample the dynamics and composition of the rich plasma environment. The reprogrammed IPS diagnostic sensors may enhance the plasma physics return and perhaps allow measurements of dust impacts.

The planned encounter of DS1 will be on Borrelly's thirteenth recorded apparition. (Unfavorable orbital conditions prevented it from being recovered on the fifth and sixth returns after its discovery.) With a period of 6.9 years, this Jupiter-family comet has been extensively studied by many investigators; it is moderately active with a well-defined coma and tails. A'Hearn *et al.*<sup>7</sup> have identified Borrelly as a member of a compositional class of comets depleted in C-chain molecules but not in NH (all with respect to OH).

Detailed encounter plans have not yet been developed, but it is expected that the MICAS infrared imaging spectrometer will yield diagnostic data between 1.2  $\mu\text{m}$  and 2.8  $\mu\text{m}$  on the volatiles and other species exposed on the nucleus' surface. In addition, panchromatic images will be used to map the three-dimensional form of the nucleus and near-nucleus jets and perhaps other discrete structures in the coma.

PEPE will measure the flux of cometary ions and electrons from 8 eV to 33 keV as a function of energy and angle, and ion composition from 1 to 140 amu/e. The results of these measurements will provide the velocity distributions and basic plasma parameters (density, velocity, and temperature) of ions and electrons plus the composition of the cometary ions.

In addition to the direct scientific gain, the encounter is expected to help serve as a useful engineering precursor to upcoming comet missions.

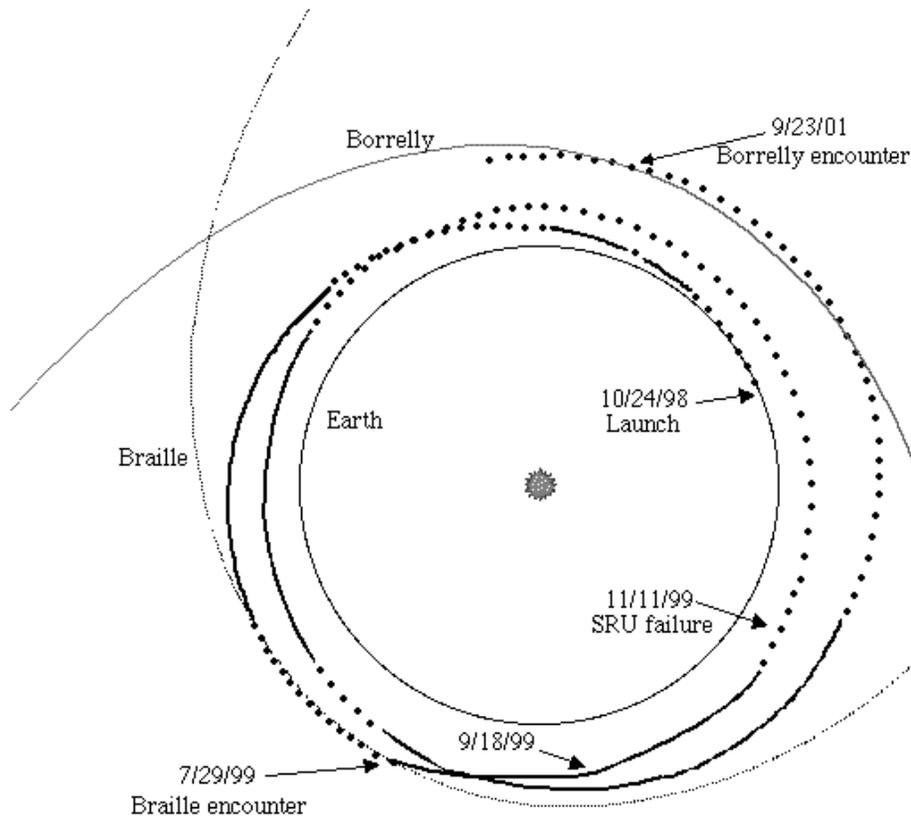


Figure 1. DS1 trajectory for the primary mission (through September 18, 1999) and extended mission. The dotted portion is for ballistic coast; the solid portion indicates the IPS is used for deterministic thrusting. As explained in the text, the IPS is on (in some cases at very low throttle levels) for most of the mission after the resumption of thrusting following the SRU failure.

### CONCLUSION

Many future missions that otherwise would have been unaffordable or even impossible now may be undertaken, with the large cost and high risk of using attractive but unproven technologies being substantially reduced because of the successful results of DS1's testing. Although the primary mission's only requirements were to assess the payload of high-risk technologies, the extended mission offers an opportunity to go beyond those objectives and to return important science data from comet Borrelly. The failure of the SRU early in the extended mission made it unlikely that DS1 would accomplish anything further of value, but the ambitious and successful recovery has restored the spacecraft and returned the mission to smooth operations. While a great deal of work and further risks

remain before the encounter, the potential returns make the DS1 extended mission attractive indeed.

### ACKNOWLEDGMENTS

The members of the DS1 mission operations and anomaly recovery team are gratefully acknowledged for their meritorious and even heroic work, upon which this overview is based. The support of JPL and NASA HQ management during this very difficult recovery effort is also greatly appreciated.

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

## REFERENCES

1. Rayman, Marc D. and David H. Lehman, "NASA's First New Millennium Deep-Space Technology Validation Flight," Second IAA International Conference on Low-Cost Planetary Missions, Laurel, MD, April 16-19, 1996, IAA-L-0502.
2. Rayman, Marc D., Philip Varghese, David H. Lehman, Leslie L. Livesay, "Results from the Deep Space 1 Technology Validation Mission," *Acta Astronautica* **47**, p. 475 (2000).
3. Deep Space 1 Technology Validation Reports, JPL Publication 00-10 (Jet Propulsion Laboratory, Pasadena, CA, to be published in 2000).
4. Rayman, Marc D. and David H. Lehman, "Deep Space One: NASA's First Deep-Space Technology Validation Mission," *Acta Astronautica* **41**, p. 289 (1997).
5. Soderblom, L. A., "DS1 Micas Full-Disk-Rotation Spectra Of Mars (1.2-2.8 Microns)," Submitted to 32nd Annual AAS Division for Planetary Sciences Meeting, October 23-27, 2000, Pasadena, CA.
6. Rayman, Marc D., Pamela A. Chadbourne, Jeffery S. Culwell, and Steven N. Williams, "Mission Design for Deep Space 1: A Low-Thrust Technology Validation Mission," Third IAA International Conference on Low-Cost Planetary Missions, Pasadena, CA, April 27 - May 1, 1998, IAA-L98-0502.
7. A'Hearn, Michael F., Robert L. Millis, David G. Schleicher, David J. Osip, and Peter V. Birch, "The Ensemble Properties of Comets: Results from Narrowband Photometry of 85 Comets, 1976-1992," *Icarus* **118**, p. 223 (1995).